

Abstract of the Invention

InsA²₅

A digital data processing system with improved network access to files in a storage system has a plurality of digital data processing nodes and a storage device. First or "client" nodes locally cache meta data such as maps to file data stored on the storage device and directly mount the file system using their cached data. A second "server" or MDC node connects to the storage device and mediates access to files, issuing file change notifications over the network for changed files or meta data. The second node preferably implements a SAN shared access storage system, and initially provides the maps for requested files, from which each of the first nodes builds their own

10 cache in a natural way. In further embodiments, the client nodes may register for access to particular file areas or sub-directories, and may then receive change notifications only for their registered access areas or portions of the data tree; alternatively, the second node may determine scope of necessary notifications from the set of previously-accessed files. A file application on the client inspects change lists to determine whether to directly access (or use previously accessed file

15 data) using the cached information, or whether a desired file must be accessed by requesting new maps from the second node through a network file sharing protocol. The system operates on the premise that large parts of the file system remain unchanged for a substantial time. Thus, rather than arbitrating each request for file access, the file notification system publishes and updates one or more lists of file changes. These may identify portions of the file directory cached at a client node that are

20 either off-limits, or require specific determinations as to their continued usability. The local file application implements a two-stage access protocol that allows direct mounting of the file system by all or a subset of client nodes for both files and meta data, and extends the time during which file accesses by cached data may be effected. The notification system may run on a server node configured to implement shared storage by interfacing with and coordinating file access, write and

25 block re-allocation operations of a native file management system to maintain file integrity, coherence or security, and may apply the change list to prevent local file applications from accessing restricted data as well as preventing client node use of changed file data or outdated data maps.

908645.1